

TITLE OF THE INVENTION

BROADBAND ACCESS FOR VIRTUAL PRIVATE NETWORKS

INVENTORS

Keith ALLEN

Weijing CHEN

P23664.S03

BROADBAND ACCESS FOR VIRTUAL PRIVATE NETWORKS

BACKGROUND OF THE INVENTION

1. Field of the Invention

[0001] The present invention relates to the field of telecommunications. More particularly, the present invention relates to using an Internet protocol (IP) network of a telecommunications service provider to provide virtual private network (VPN) functionality among local area networks (LANs).

2. Acronyms

[0002] The written description provided herein contains acronyms which refer to various telecommunications services, components and techniques, as well as features relating to the present invention. Although some of these acronyms are known, use of these acronyms is not strictly standardized in the art. For purposes of the written description herein, the acronyms are defined as follows:

Address Resolution Protocol (ARP)

Asynchronous Transfer Mode (ATM)

Digital Subscriber Line (DSL)

Digital Subscriber Line Access Multiplexer (DSLAM)

Internet Protocol (IP)

Internet Protocol Version 4 (IPv4)

Internet Protocol Version 6 (IPv6)

Internet Service Provider (ISP)

Local Area Network (LAN)

Media Access Control (MAC)

Multi-Protocol Label Switching (MPLS)

Point-to-Point Protocol (PPP)

Personal Digital Assistant (PDA)

Request for Comment (RFC)

Telecommunications Service Provider (TSP)

Transmission Control Protocol (TCP)

User Datagram Protocol (UDP)

Virtual Local Area Network (VLAN)

Virtual Private Network (VPN)

Wide Area Network (WAN)

3. Background and Material Information

[0003] A virtual private network (VPN) is a non-public network that runs over a shared network infrastructure, such as the public Internet. A VPN enables interconnection among distinct networks, including local area networks (LANs), and other end-systems over a wide area network (WAN). The VPN provides security to the extent that it recognizes and transports only data associated with end-systems that are part of the network.

[0004] A virtual local area network (VLAN) provides logical grouping and networking of various customer end-systems, such as work stations, user devices, private networks, and the like, as though they are grouped on the same physical LAN. In other words, the VLAN associates end-systems based on criteria other than the physical location of the end-systems. For example, the VLAN can provide network services to a customer having multiple geographic locations, or to a department within a customer organization remotely located throughout a campus environment.

[0005] A VLAN VPN implemented over a WAN is provided by a telecommunications service provider (TSP) to interconnect the LANs as if they were one. Typically, TSPs rely on network connections to provide customers VPN service,

including multi-protocol label switching (MPLS) paths, asynchronous transfer mode (ATM) circuits and point-to-point protocol (PPP) connections. However, connection oriented implementations inhibit efficient execution of VPN services and have limited scalability, as well as flexibility.

[0006] The present invention overcomes the problems associated with the prior art, as described below.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] The present invention is further described in the detailed description that follows, by reference to the noted drawings by way of non-limiting examples of embodiments of the present invention, in which like reference numerals represent similar parts throughout several views of the drawings, and in which:

Fig. 1 is a diagram showing an exemplary network architecture, according to an aspect of the present invention; and

Fig. 2 is a flow diagram illustrating IP broadband access from an originating LAN, according to an aspect of the present invention.

DETAILED DESCRIPTION OF EMBODIMENTS

[0008] The present invention incorporates an Internet protocol, version 6 (IPv6) network to provide a connectionless approach to implementing VLAN VPNs among various LANs. Generally, the invention enables broadband access to the IPv6 network of a telecommunications service provider (TSP) by emulating layer two (*e.g.*, Ethernet) functionality while performing layer three (*e.g.*, IP) routing through a WAN. For example, a user operating within an originating LAN of a VPN may access the IPv6 TSP network from either a IPv4 or an IPv6 originating device by encapsulating each frame in an IPv6 packet, together with a unique VPN identification number, at an ingress line interface of the TSP network. The LAN

frames include, for example, Ethernet frames, media access control (MAC) frames, other layer two frames and the like. The IPv6 packet is routed through the TSP network to an egress line interface associated with the destination LAN within the same VPN. The egress line interface verifies the VPN identification number and, upon verification, decapsulates the IPv6 packet and transmits the LAN frame to the destination LAN. Use of the IPv6 TSP network is a scalable and more manageable alternative to current MPLS networks, for example.

[0009] In view of the above, the present invention through one or more of its various aspects and/or embodiments is presented to accomplish one or more objectives and advantages, such as those noted below.

[0010] An aspect of the present invention provides a method for sending data through a provider network from an originating network to a destination network in a VPN. The method includes encapsulating data link layer data from the originating network in a network layer packet and determining whether a data link layer address of a destination device in the destination network is mapped to a network layer address of an egress line interface in the provider network. When the destination device address is not mapped to the egress line interface address, the network layer packet is broadcast to a multicast address associated with the VPN. When the destination device address is mapped to the egress line interface address, based on a previous transmission from the destination device, the network layer packet is unicast to the egress line interface address.

[0011] A VPN identification number corresponding to the VPN may be added to the network layer packet. The VPN identification number is verified after the egress line interface receives the network layer packet. The data layer link data is

decapsulated from the network layer packet only when the VPN identification number is verified.

[0012] Another aspect of the present invention provides a method for providing broadband access to a VPN, which includes multiple LANs configured to interface with an IPv6 service provider network through broadband access links. The method includes encapsulating a LAN frame from an originating LAN of the VPN in an IPv6 packet of the service provider network; adding a VPN identification number corresponding to the VPN to the IPv6 packet; and routing the IPv6 packet through the service provider network. The LAN frame is decapsulated when the VPN identification number is verified. The decapsulated LAN frame is transmitted to the destination LAN. The IPv6 packet is discarded when the VPN identification number is not verified. The IPv6 packet includes an IPv6 address of an ingress line interface, which receives the LAN frame, as a source address and an IPv6 address of an egress line interface, to which the IPv6 packet is routed for verification, as a destination address.

[0013] The IPv6 packet may include the VPN identification number in an optional header extension. For example, the VPN identification number may be included in a multiple of four octets of the optional header extension. The optional header extension may further identify a destination option type, in which case the method further includes discarding the IPv6 packet when the egress line interface does not recognize the destination option type in the optional header extension. The optional header extension may further identify a VPN hop number, which indicates a number of line interfaces that transmit the IPv6 packet.

[0014] It is determined whether an address of a destination device in the destination LAN is mapped to the egress line interface. When the address is not

mapped to the egress line interface, the IPv6 packet is broadcast to a multicast address associated with the VPN and the IPv6 packet is received at the egress line interface based on the multicast address. An address of the egress line interface is mapped to the address of the destination device, based on address information received by the ingress line interface in a transmission from the destination device. Subsequent IPv6 packets are then transmitted to the destination device using a unicast address of the egress line interface based on the mapping.

[0015] Another aspect of the present invention provides a system for providing broadband access to a VPN, which includes multiple LANs configured to interface with an IPv6 service provider network. The system includes multiple interface devices in the service provider network. Each interface device includes at least one line interface, each of which is connectable to at least one of the LANs through a broadband access link. A first interface device receives a LAN frame from a first LAN at an ingress line interface corresponding to the first LAN, encapsulates the LAN frame in an IPv6 packet, and adds a VPN identification number corresponding to the VPN to the IPv6 packet. The LAN frame is directed to a second LAN. A second interface device receives the IPv6 packet at an egress line interface corresponding to the second LAN, verifies the VPN identification number, decapsulates the LAN frame when the VPN identification number is verified, and transmits the LAN frame to the second LAN. The second interface device discards the IPv6 packet when it is not able to verify the VPN identification number.

[0016] The IPv6 packet includes the VPN identification number in an optional header extension. The first interface device may further include an ingress virtual bridge corresponding to the ingress line interface. When the ingress virtual bridge is not able to associate an address of a destination device in the second LAN with an

address of the egress line interface of the second interface device, the first interface device broadcasts the IPv6 packet to a multicast address associated with the VPN. The second interface device receives the IPv6 packet at the egress line interface based on the multicast address. The second interface device may further include an egress virtual bridge corresponding to the egress line interface. The egress virtual bridge then maps an address of an originating device in the first LAN with the address of the ingress line interface after the second interface device receives the broadcast IPv6 packet. The second interface device is able to unicast subsequent IPv6 packets, directed to the originating device, to the address of the ingress line interface based on the mapping.

[0017] Yet another aspect of the present invention provides a method of providing broadband access for a customer in a VPN, including multiple LANs interfacing with at least one TSP network. Each TSP network includes multiple interfaces corresponding to the multiple LANs. The method includes assigning a unique VPN identification number to the customer and assigning a common multicast address to the interfaces and a unique unicast address to each of the interfaces. Data from an originating LAN is received, the data being directed to a destination device in a destination LAN. The originating LAN corresponds to an ingress interface of the interfaces and the destination LAN corresponds to an egress interface of the interfaces. When an address of the destination device address is not mapped to the destination LAN, the data is encapsulated in a multicast packet, having the unique address of the ingress interface as a source address and the multicast address as a destination address. The encapsulated data is transmitted to all interfaces corresponding to the LANs based on the multicast address. The frame is decapsulated only at the egress interface, which forwards the frame to the destination device.

[0018] The destination device address may be mapped to the IPv6 address of the egress interface based on address information previously received by the ingress interface from the destination device. When the destination device address is mapped to the egress interface, the data is encapsulated in a unicast packet having the unique IPv6 address of the ingress interface as the source address and the unique address of the egress interface as the destination address. The encapsulated frame is then transmitted only to the egress interface, based on the destination address in the unicast packet. The VPN identification number may be entered in the multicast packet and/or the unicast packet. The reading the VPN identification is read at the egress interface to verify that the received packet is associated with the VPN.

[0019] The various aspects and embodiments of the present invention are described in detail below.

[0020] The present invention is directed to enabling TSPs to provide IP broadband network services over packet switched data networks having expanded capacity, such as IPv6 networks. IPv6 networks are capable of serving mass-market IP broadband access subscribers, as well as accommodating business customers to manage their data communication services in-house. VLAN VPN is one service that enables the broadband access and management control desired.

[0021] Fig. 1 is a diagram depicting an exemplary network infrastructure supporting the present invention. Fig. 1, in particular, depicts a VLAN VPN of a customer, which includes multiple LANs 10, 40 and 50. In an embodiment of the invention, each LAN is an Ethernet LAN having user end-systems 14, 44 and 54 that interface with a TSP core network 20 through customer edge devices 11, 41 and 51, respectively. However, each of the LANs 10, 40 and 50 may be an IP network or other data network without departing from the spirit and scope of the present

invention. As discussed above, the TSP network 20 is an IPv6 network, which allows static allocation of IPv6 addresses to uniquely identify the customers, due to the large address space of the IPv6 format. The TSP IPv6 network 20 is essentially configured such that it appears to be a wide-area LAN to the VLAN VPN customer.

[0022] The LANs 10, 40 and 50, depicted in the exemplary embodiment of the invention, provide layer two (*i.e.*, data link layer) network functionality. Layer two functionality generally includes handling physical and logical connections to the Ethernet or other LAN frame (or packet) destinations. The user devices, such as customer devices 14, 44 and 54, are addressed and identified using unique MAC addresses corresponding to each device. The data link layer protocol specifies the MAC address of each frame's source and destination.

[0023] The TSP network 20, depicted in the exemplary embodiment of the invention, provides layer three (*i.e.*, network layer) functionality. The network layer routes packets (or datagrams) from one network to another. Internet protocol, in particular, identifies each IP device with a unique IP address, including, for example, the edge devices 22, 23 and 24 and/or the various line interfaces 22a, 23a and 24a, discussed below. In the depicted embodiment of the present invention, each IP address is an IPv6 address allocated by the TSP network 20. The IP packets are routed through the TSP network 20 in accordance with IP (*e.g.*, IPv6), while the layer four (*i.e.*, transport layer) protocols, such as transmission control protocol (TCP) and user datagram protocol (UDP) for transmission control, continue to be supported with no change.

[0024] The originating customer device 14 is depicted as part of the originating LAN 10. Each of the LANs 10, 40 and 50 may include any number and type of IP compatible networked devices, including, for example, a personal computer, a laptop

computer, a personal digital assistant (PDA), a voice over IP telephone, or the like. The originating customer device 14 communicates through the LAN 10 with the customer edge device 11. The customer edge device 11 likewise is any interface device capable of communicating with the TSP network 20, such as an Ethernet switch or an IP router with an Ethernet interface, depending on the type of customer network.

[0025] The customer edge device 11 accesses an ingress router 22 of the TSP network 20 over a broadband access link 12. In an embodiment of the invention, the broadband access link 12 is a digital subscriber line (DSL), and therefore includes a DSL access multiplexer (DSLAM), an asynchronous transfer mode (ATM) edge switch and an interworking function device (not pictured), for example. Alternatively, the broadband access link 12 may include digital cable, T-1, digital signal-level 3 (DS-3) or optical carrier-level 3 (OC-3) interfaces, or an Ethernet. The broadband access links 12, 42 and 52 of the various LANs in the VPN do not need to be the same type.

[0026] The TSP network 20 includes multiple edge devices (*i.e.*, gateways) 22, 23 and 24, which may be IPv6 routers or switches having Ethernet bridging functionality, for example. Each edge device 22, 23 and 24 has multiple line interfaces, which correspond to the various customer LANs and interface the edge devices with the customer LANs, including depicted line interfaces 22a, 23a and 24a. It is understood that each edge device 22, 23 and 24 is configured with multiple line interfaces, even though Fig. 1 depicts only one corresponding line interface 22a, 23a and 24a for each. The multiple line interfaces enable each edge device 22, 23 and 24 to simultaneously service multiple VPNs and LANs.

[0027] Typically, each customer device 14 , 44 and 54 corresponds to a single line interface (*e.g.*, line interfaces 22a, 23a and 24a), although alternative embodiments of the invention include multiple interfaces for a single customer device 14, 44 and 54. As discussed below, each line interface 22a, 23a and 24a is associated with a virtual learning bridge, which learns and caches mapping of customer devices 14, 44 and 54 with their associated line interfaces 22a, 23a and 24a. The TSP network 20 is thus able to more efficiently route communications among the LANs in the VPN over time.

[0028] The exemplary network architecture of Fig. 1 depicts an ingress edge device 22, which interfaces with the broadband access link 12 of the originating LAN 10 through an ingress line interface 22a. In addition to the ingress edge device 22, Fig. 1 depicts two egress edge devices 23 and 24. Egress edge device 23 interfaces with the broadband access link 42 of the destination LAN 40 through an egress line interface 23a, and the egress edge device 24 interfaces with the broadband access link 52 of the destination LAN 50 through an egress line interface 24a. Each of the LANs 40 and 50 respectively include at least one destination device 44 and 54 of the customer, which may be any of the types of devices described above with respect to the originating customer device 14 of LAN 10. It is understood that all of the exemplary customer devices and corresponding edge devices depicted in Fig. 1 are capable of sending and receiving data through the VPN, and are described as originating or destination devices only to facilitate description of the various embodiments of the invention.

[0029] IPv6 packets are routed through the TSP network 20 to one of the egress edge devices 23 or 24, depending on the destination address of each IPv6 packet sent from the ingress edge device 22. For example, an Ethernet frame from the customer

device 14 destined for the destination device 44 in LAN 40 is encapsulated in an IPv6 packet having the IP address of the egress line interface 23a as the destination address. As described in detail with respect to Fig. 2, when the egress line interface associated with the destination device and/or the destination LAN is not known, the ingress edge device 22 multicasts the IPv6 packet to the other edge devices serving the VLAN VPN so that all potential egress edge devices, including the egress edge devices 23 and 24, receive the packet. Only the egress edge device having the egress line interface that services the destination LAN (*e.g.*, the egress edge device 23 and associated egress line interface 23a) responds to the broadcast IPv6 packet, enabling unicast transmission of subsequent packets from the ingress edge device 22. In alternative embodiments of the invention, the broadcast comprises a modified address resolution protocol (ARP) message and the encapsulation and decapsulation is performed by the customer edge devices 11 and 41. The addressing and transmitting processes is described in detail with respect to Fig. 2, below.

[0030] In an embodiment of the invention, the customer subscribes to a VLAN VPN service having IP broadband connectivity with the TSP network 20 from multiple sites or locations. The customer's VLAN VPN is assigned a unique VPN identification number, which is four or more bytes, for example. In alternative embodiments of the invention, the unique VPN identification number is an IPv6 address prefix under control of the TSP or an identification number allocated from a dedicated, separate address space of the TSP network 20. The VPN identification number assigned to the customer is included with the VPN interface configuration. The VPN identification number distinguishes data packets associated with the customer's VLAN VPN from other VLAN VPNs supported by the TSP network 20.

[0031] Each of the line interfaces 22a, 23a and 24a depicted in the TSP network 20 is allocated unique IPv6 address (TSP-IPv6) from an IPv6 address block of the TSP network 20. As explained above, the line interfaces 22a, 23a and 24a and associated TSP-IPv6 addresses correspond to particular LANs of the customer's VPN. The line interfaces 22a, 23a and 24a are also assigned a VLAN VPN specific IPv6 multicast address (TSP-MIPv6) from the TSP's IPv6 address block, associated with the customer's VPN. A single multicast address may be used to multicast a packet to all of the interfaces serving a VLAN VPN. In an embodiment of the invention, the TSP-IPv6 and TSP-MIPv6 addresses are allocated or assigned to the line interfaces 22a, 23a and 24a manually, although any effective form of allocation or assignment may be incorporated without departing from the spirit and scope of the present invention. Generally, using the IPv6 and the MIPv6 addresses, the TSP network 20 is able to effectively provide layer two (*e.g.*, Ethernet) capabilities to link customer edge devices 11, 41 and 51, while the customer performs its own layer three (*e.g.*, IP) network administration.

[0032] As stated above, the ingress edge device 22 has virtual learning bridges corresponding to the line interfaces associated with every VLAN VPN that it serves, including the line interface 22a. When the virtual learning bridge receives an Ethernet frame from the originating VLAN 10, for example, it learns and caches identification information, such as an Ethernet MAC address and/or the LAN identification number (*e.g.*, the 802.1q tag) of the originating customer device 14 from which the frame is sent. Therefore, the line interface 22a knows precisely where to forward frames that it subsequently receives, *e.g.*, from the other line interfaces 23a and 24a, destined for the MAC address and/or the VLAN identification number of the originating customer device 14.

[0033] Similarly, the virtual learning bridge learns information that enables the line interface 22a to efficiently forward LAN frames that it receives from the originating customer device 14 to various destination devices, such as the destination customer devices 44 and 54. For example, referring to Fig. 2, the ingress line interface 22a

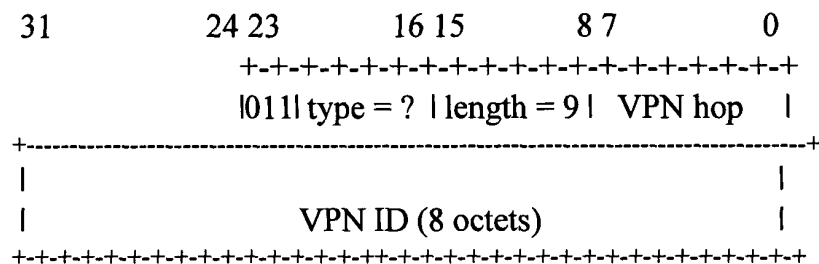
[0034] receives a LAN frame from the originating customer device 14 over the broadband access link 12 at step s210. At step s212, the ingress line interface 22a determines whether the destination address (*e.g.*, the address of the destination customer device 44) is already mapped to an egress line interface.

[0035] When mapping of the destination address to an egress line interface does not exist, for example, when the ingress line interface 22a has no record of transmitting a packet from the originating customer device 14 to the destination customer device 44 (or from the originating LAN 10 to the destination LAN 40), the ingress line interface 22a encapsulates the LAN frame in a multicast IPv6 packet at step s214. The multicast IPv6 packet includes the TSP IPv6 address of the ingress line interface 22a as the source address and the TSP-MIPv6 multicast address of the customer's VLAN VPN as the destination multicast address.

[0036] At step s216, the previously assigned VPN identification number is added to the header of the IPv6 packet in order to provide security for the VPN. Use of the VPN identification number prevents unsecured or unauthorized LAN frames from being delivered to VPN customers at the egress line interfaces 23a and 24a. Use of the VPN identification number is an improvement over security measures implemented in conventional VPNs, which typically require layered connections internal to the network, and a control plane consisting of virtual routers configured to exchange routing information. In an alternative embodiment, the customer edge device 11 adds the VPN identification number to the IPv6 packet header, and the

ingress line interface 22a confirms it. The IPv6 packet, including the VPN identification number is transmitted through the TSP network 20 to the egress line interface 23a in the egress device 23 at step s220.

[0037] In an embodiment of the invention, an extension of the current IPv6 optional header is used to encapsulate the VPN identification number. The IPv6 header is implemented, for example, in accordance with RFC 2460, “Internet Protocol, Version 6 (IPv6) Specification” (December 1998), the content of which is expressly incorporated by reference herein in its entirety. The VPN identification number marks an IPv6 packet as belonging to a particular VLAN VPN. The VPN identification number header is a specific option of the more generic destination options header (*e.g.*, header type 60) of the IPv6 protocol. An exemplary format of the optional header extension, including the VPN identification number header is as follows:



[0038] The first three bits of the first octet are 011, as shown above. The remaining five bits comprise the destination option type number. The value of 011 indicates that nodes not recognizing this option type should discard the packet and that the option data (*i.e.*, the VPN hop count) may change en route. Discarding the packet ensures that any packet delivered to a node not capable of processing VPN headers will not be inadvertently delivered to a site outside of the VPN. The VPN hop count is an eight bit unsigned integer, which is incremented by one by each

peering edge device in the TSP network 20 that forwards the packet. The VPN identification number is a four (or multiple of four) octet identifier associated with each VPN.

[0039] At step s218, the multicast IPv6 packet, including the encapsulated LAN frame and the VPN identification number, is then broadcast through the TSP network 20, resulting in the IPv6 packet being received by every potential egress line interface associated with a broadband access of the customer's VLAN VPN, including, for example, line interfaces 23a and 24a, at step s220. Only the line interface 23a which corresponds to the destination LAN 40 and/or the destination customer device 44 proceeds with the remaining steps of Fig. 2, for example, based on the MAC address of the destination customer device 44 to which the LAN frame was initially addressed.

[0040] When the mapping exists, as determined at step s212, the ingress line interface 22a encapsulates the LAN frame in a unicast IPv6 packet at step s232. The IPv6 packet has the TSP IPv6 address of the egress line interface (*e.g.*, the line interface 23a) as the destination address and the TSP IPv6 address of the ingress line interface 22a as the source address. At step s234, the previously assigned VPN identification number is added to the header of the IPv6 packet in order to provide security for the VPN, as described above with respect to multicasting IPv6 packets. The unicast IPv6 packet, including the encapsulated LAN frame and the VPN identification number, is then transmitted through the TSP network 20 using the TSP IPv6 address of the egress line interface 23a, which receives the unicast IPv6 packet at step s220.

[0041] At step s222, the virtual learning bridge of the egress line interface 23a authenticates the IPv6 packet. For example, the egress line interface 23a first

determines whether the VPN identification number of the received IPv6 packet matches the assigned customer VPN identification number. Any IPv6 packets that do not include a matching VPN identification number are discarded at step s238. In an embodiment of the invention, the destination customer edge device 41 determines whether the VPN identification number of the received IPv6 packet matches the assigned VPN identification number and discards unauthorized packets, accordingly. In another embodiment, authorization of the VPN identification numbers can be disabled in the TSP network 20 and/or the LAN 40 to enable interworking among a greater number of VPNs.

[0042] When the VPN identification number of the IPv6 packet matches the customer VPN identification number, the virtual learning bridge of the egress line interface 23a decapsulates the IPv6 packet and extracts the LAN frame at step s224. The LAN frame is forwarded to the destination LAN 40, through the broadband access link 42 at step s226.

[0043] Meanwhile, at step s228, the virtual learning bridge of the egress line interface 23a learns and caches the mapping of identification information, such as an Ethernet MAC address and/or the VLAN identification number of the originating customer device 14, to the TSP-IPv6 address of the ingress line interface 22a, from which the frame was sent. Therefore, when the egress line interface 23a receives subsequent LAN frames from the customer device 44 and/or the LAN 40, destined for the customer device 14 and/or the LAN 10, the egress line interface 23a merely encapsulates the LAN frame in an IPv6 packet and unicasts the IPv6 packet to the ingress line interface 22a, using the mapping.

[0044] Because the majority of Ethernet data traffic, for example, is bi-directional, eventually all of the line interfaces 22a, 23a and 24a will learn and cache the mapping

of the identification information of all active customer devices 14, 44 and 54 in all of the LANs 10, 40 and 50 to the TSP IPv6 addresses of the corresponding line interfaces 22a, 23a and 24a. Accordingly, the line interfaces will be able to encapsulate the LAN frames in unicast IPv6 packets containing the specific TSP-IPv6 address of the desired line interface as the destination address, instead of a multicast IPv6 packet having the customer IPv6 multicast address as the destination address. In other words, as the mapping among line interfaces 22a, 23a and 24a and customer devices 14, 44 and 54 is learned and cached, the process increasingly follows the unicast steps s232, s234 and s236 of Fig. 2, as opposed to the multicast steps s214, s216 and s218, significantly increasing the efficiency of the TSP network 20.

[0045] An alternative embodiment of the present invention enables interworking among different VLAN VPNs (*i.e.*, extra-net VPNs). The line interfaces 22a, 23a and 24a are configured to recognize and authenticate multiple, previously assigned VPN identification numbers, corresponding to the interworking VPNs, instead of a single VPN identification number corresponding to one customer. Accordingly, any IPv6 packet that arrives at the line interfaces 22a, 23a and 24a having a VPN identification number matching any of the VPN identification numbers on the list is authenticated and forwarded to the appropriate LAN 10, 40 and 50 and/or customer device 14, 44 and 54. For example, the allowed VPN identification number list may include all of the VPN identification numbers of pre-arranged peering business customers.

[0046] Although the invention has been described with reference to several exemplary embodiments, it is understood that the words that have been used are words of description and illustration, rather than words of limitation. Changes may be made within the purview of the appended claims, as presently stated and as amended, without departing from the scope and spirit of the invention in its aspects.

Although the invention has been described with reference to particular means, materials and embodiments, the invention is not intended to be limited to the particulars disclosed; rather, the invention extends to all functionally equivalent structures, methods, and uses such as are within the scope of the appended claims.

[0047] In accordance with various embodiments of the present invention, the methods described herein are intended for operation as software programs running on a computer processor. Dedicated hardware implementations including, but not limited to, application specific integrated circuits, programmable logic arrays and other hardware devices can likewise be constructed to implement the methods described herein. Furthermore, alternative software implementations including, but not limited to, distributed processing or component/object distributed processing, parallel processing, or virtual machine processing can also be constructed to implement the methods described herein.

[0048] It should also be noted that the software implementations of the present invention as described herein are optionally stored on a tangible storage medium, such as: a magnetic medium such as a disk or tape; a magneto-optical or optical medium such as a disk; or a solid state medium such as a memory card or other package that houses one or more read-only (non-volatile) memories, random access memories, or other re-writable (volatile) memories. A digital file attachment to email or other self-contained information archive or set of archives is considered a distribution medium equivalent to a tangible storage medium. Accordingly, the invention is considered to include a tangible storage medium or distribution medium, as listed herein and including art-recognized equivalents and successor media, in which the software implementations herein are stored.

[0049] Although the present specification describes components and functions implemented in the embodiments with reference to particular standards and protocols, the invention is not limited to such standards and protocols. Each of the standards for Internet and other packet-switched network transmission (e.g., IPv4, IPv6, TCP, UDP, MPLS) and public telephone networks (ATM, DSL) represent examples of the state of the art. Such standards are periodically superseded by faster or more efficient equivalents having essentially the same functions. Accordingly, replacement standards and protocols having the same functions are considered equivalents.